

AFRL-SN-WP-TR-2006-1127

**SENSOR MANAGEMENT FOR
FIGHTER APPLICATIONS**

**Stanton H. Musick
Raj P. Malhotra**



**ATR and Fusion Algorithms Branch (AFRL/SNAT)
Sensor ATR Technology Division
Sensors Directorate
Air Force Research Laboratory, Air Force Materiel Command
Wright-Patterson Air Force Base, OH 45433-7320**

JUNE 2006

Final Report for 01 October 1993 – 30 September 2004

Approved for public release; distribution is unlimited.

STINFO COPY

**SENSORS DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7320**

NOTICE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory Wright Site (AFRL/WS) Public Affairs Office (PAO) and is releasable to the National Technical Information Service (NTIS). It will be available to the general public, including foreign nationals.

PAO Case Number: AFRL/WS-06-1646, 29 Jun 2006.

THIS TECHNICAL REPORT IS APPROVED FOR PUBLICATION.

//Signature//

STANTON H. MUSICK
Project Engineer
ATR & Fusion Algorithms Branch
Sensor ATR Technology Division

//Signature//

DALE E. NELSON, Ph.D.
Chief, ATR & Fusion Algorithms Branch
Sensor ATR Technology Division
Sensors Directorate

//Signature//

PATRICK D. KEE
Lt. Col., USAF
Deputy, Sensor ATR Technology Division
Sensors Directorate

This report is published in the interest of scientific and technical information exchange and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YY) June 2006		2. REPORT TYPE Final		3. DATES COVERED (From - To) 10/01/1993 – 09/30/2004	
4. TITLE AND SUBTITLE SENSOR MANAGEMENT FOR FIGHTER APPLICATIONS				5a. CONTRACT NUMBER In-house	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 61102F	
6. AUTHOR(S) Stanton H. Musick Raj P. Malhotra				5d. PROJECT NUMBER 2304	
				5e. TASK NUMBER ES	
				5f. WORK UNIT NUMBER 01	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) ATR and Fusion Algorithms Branch (AFRL/SNAT) Sensor ATR Technology Division Sensors Directorate Air Force Research Laboratory, Air Force Materiel Command Wright-Patterson Air Force Base, OH 45433-7320				8. PERFORMING ORGANIZATION REPORT NUMBER AFRL-SN-WP-TR-2006-1127	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Sensors Directorate Air Force Research Laboratory Air Force Materiel Command Wright-Patterson AFB, OH 45433-7320				10. SPONSORING/MONITORING AGENCY ACRONYM(S) AFRL-SN-WP	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S) AFRL-SN-WP-TR-2006-1127	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES Report contains color. PAO Case Number: AFRL/WS-06-1646, 29 Jun 2006.					
14. ABSTRACT This report is a summary of research conducted from 1 October 1993 through 30 September 2004 for the project titled <i>Sensor Management for Fighter Applications</i> (SMFA). This project developed techniques for intelligently allocating the sensors onboard a modern military aircraft. It focused on information metrics for balancing the needs of detection, tracking and identification, on a probabilistic representation for assimilating sensed data in a multitarget environment, on machine learning approaches, and on important applications of these technologies. This report is the final written document for this project.					
15. SUBJECT TERMS sensor allocation, sensor resource management, joint multitarget probability, discrimination gain					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT: SAR	18. NUMBER OF PAGES 40	19a. NAME OF RESPONSIBLE PERSON (Monitor) Stanton H. Musick 19b. TELEPHONE NUMBER (Include Area Code) N/A
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			

Contents

Contents	iii
Figures.....	iv
Acronyms	v
Introduction.....	1
Staffing.....	2
Problem Background	2
A Typical Problem and Its Technical Issues.....	3
Setting the Stage	4
Mathematical Results.....	5
Discrimination Gain.....	5
Joint Multitarget Probability	7
Nonlinear Filtering.....	9
TENET	13
Applications of Nonlinear Filtering.....	13
Tracking in the presence of multipath interference	14
Association-free bias estimation	14
Tracking through radar clutter	15
A method for finding distributed objects	15
Machine Learning Results	17
Reinforcement Learning in Sensor Management	17
Virtual Associative Networks for Sensor Management.....	19
New Directions	22
References.....	28

Figures

Figure 1. Comparison of four search methods.....	6
Figure 2. Sensing effort as a function of location and time for the 2-target example.	8
Figure 3. JMP probability P_2 that there are 2 targets in the scene.	9
Figure 4. A single image at -3 dB, target at (27, 13)	11
Figure 5. Position marginal, average over 10 runs	12
Figure 6. Low SNR image tracking, average over 10 runs, true dotted, estimate solid ..	12
Figure 7. An illustration contrasting managed and non-managed tracking performance	24
Figure 8. Quantitative comparison of intelligent and naïve sensor management.	25
Figure 9. Execution time for myopic sensor management	26

Acronyms

ADI	alternating direction implicit
ATR	automatic target recognition
DARPA	Defense Advanced Research Projects Agency
DP	dynamic programming
FPE	Fokker-Planck equation
GD	General Dynamics (Advanced Information Systems Division)
GMTI	ground moving target indications
HRR	high range resolution
ISP	Integrated Sensing and Processing
ISR	intelligence, reconnaissance and surveillance
JMP	joint multitarget probability
JMPD	joint multitarget probability density
LQG	linear quadratic Gaussian
MDP	Markov decision process
ML	machine learning
NLF	nonlinear filtering
ODE	ordinary differential equation
PDE	partial differential equation
SMFA	Sensor Management for Fighter Applications
SCNR	signal-to-clutter-plus-noise ratio
SNR	signal-to-noise ratio
TBD	track-before-detect
TDL	Temporal Difference Learning
TENET	TEchniques for the Nonlinear Estimation of Tracks
UAV	unmanned aerial vehicle
VAN	virtual associative network

INTRODUCTION

This report is a summary of research conducted from 1 October 1993 through 30 September 2004 for the project titled *Sensor Management for Fighter Applications* (SMFA). Two United States Air Force organizations sponsored this research:

- The Mathematics and Information Sciences Directorate of the Air Force Office of Scientific Research (AFOSR/NM), and
- The Sensors Directorate of the Air Force Research Laboratory (AFRL/SN).

This report is the final written document for Project SMFA, which carried the number designation of Task 2304 ES.

In a modern aircraft, *data fusion* is the process by which the target environment is measured by sensors, and data from sensing actions are combined into estimates, reasoned over, and presented to the pilot. Determining which data to measure and when to take those measurements is critical to achieving effective data fusion. But the need for data depends on uncertain, interrelated, and dynamic factors. This fact has pushed the activity of planning and scheduling data-sensing beyond the ability of the pilot, and has led researchers to study structured decision-aiding systems called *sensor managers*.

A sensor manager must consider questions such as where each sensor should point, what mode of sensing should be used, and how the sensors should be sequenced in time. Effective *sensor management* produces early target detection, accurate target track, and clear-cut identification. The objective of this research was to develop automated methods for the intelligent allocation of agile airborne sensors in a real-time environment comprised of targets that can be observed by those sensors.

Acceptable sensor allocation methods must be able to handle heterogeneous sensor types, targets that are moving and at rest, finite sensing assets, imperfect sensed data, and situation uncertainties. A sensor manager performs its work by identifying needs, by determining which available sensors can satisfy those needs, by prioritizing potential associations of sensors to needs, by scheduling the best sensing options, and by continually adapting to the changes of a dynamic environment.

The systems-level character of managing sensors caused the domain of effort in this project to be quite broad, extending across architectural concepts, planning techniques, policies for guiding sensing actions, scheduling considerations, and mathematical frameworks for fusing data. Adaptively selecting the appropriate sensing actions under resource and operational constraints is fundamentally a problem in mathematical optimization, with connections to operations research, decision theory, stochastic estimation, information theory, machine intelligence, and data fusion.

This report is a summary, not a retelling. That is to say, this report describes the problems that were addressed, briefly summarizes the progress that was made in their solution, identifies more advanced work for which our results were foundational, and provides a full list of references that document this project. In so doing, we hope to assemble a re-

cord that is a starting point for other investigators. We will not redevelop any model, re-explain any method, or rehash in detail any experiment, and the few results we do present are merely illustrative. We are confident that those who need to dig deeper in sensor management will be able to find their way among the extensive collection of papers that we wrote for the open literature, papers that are listed here in the References section.

STAFFING

Dr. Jon Sjogren of AFOSR/NM was the manager of this 11-year project. The following individuals were its primary contributors:

- Mr. Stanton H. Musick, principal investigator, Automatic Target Recognition and Fusion Algorithms Branch of the Sensors Directorate of the Air Force Research Laboratory (AFRL/SNAT) at Wright-Patterson AFB, Ohio;
- Mr. Raj P. Malhotra, co-investigator, also of AFRL/SNAT;
- Dr. Keith Kastella, primary industry collaborator, Unisys (and its successors) in Eagan, Minnesota, and then General Dynamics in Ann Arbor, Michigan;
- Dr. Yan M. Yufik, Institute of Medical Cybernetics, Potomac, Maryland.

These five individuals also made substantial contributions:

- Dr. Wayne Schmaedeke, Unisys in Eagan, Minnesota;
- Dr. Christopher Kreucher, General Dynamics and University of Michigan, both in Ann Arbor, Michigan;
- Dr. Milton Cone, Embry-Riddle Aeronautical University in Prescott, Arizona and AFOSR Summer Faculty at Wright-Patterson AFB, Ohio;
- Mr. John Greenewald, General Dynamics and Nonlinear Vision in Dayton, Ohio.

Mr. Musick managed this project, conducted research for it, and published results over its entire 11-year history. Mr. Malhotra and Dr. Kastella participated for the first eight years 1993-2001, Dr. Cone for the five summers 1993-1997, Dr. Kreucher in the 1999-2001 time frame, and Mr. Greenewald for the period 2000-2004. From time to time, other individuals also contributed products to this work, as can be seen by the author names on various papers cited in the References section.

PROBLEM BACKGROUND

To better appreciate the sensor management problem and the challenges it presents, it is useful to look back at the situation that existed in the early 1990s when rigorous foundational work was just beginning. Here is a list (not exhaustive) of issues, assumptions and conditions that predominated at that time:

- operational sensing schedules were largely fixed, not adaptive
- in planning, sensor behavior was often assumed to be deterministic
- the possibility for target motion was sometimes ignored
- the estimation approach made unwarranted assumptions, e.g. independent targets
- single-objective optimization predominated, e.g. either identification or tracking, but not both

- often single sensor/single target reasoning
- myopic solutions

At the beginning of this project, the way forward in each of these areas was an open question. For example, how to move from fixed to dynamically adaptive sensor schedules, how to portray stochastic target motion in a manner consistent with the structure of a sensor allocation system, how to strike a reasonable balance between the objectives of search, track and identification, how to weigh the value of measurements well into the future, how best to represent the generalized multitarget problem, and so forth, were all unresolved matters.

Prior to the 1990s, designers typically approached sensor management problems in an ad hoc manner, often utilizing rule-based methods, or some combination of rules and procedures that was optimized for a subset of functions. Such solution approaches are subject to many faults. They can be brittle meaning that small deviations in scenario assumptions erodes their effectiveness disproportionately; they are usually untrustworthy because they are not guided by an underlying theory that directs their development and allows comparison to a performance bound; and they are often difficult to implement and maintain because of their specialized nature.

A Typical Problem and Its Technical Issues

To place these issues in context, consider the problem of detecting, tracking, identifying and intercepting a collection of airborne targets by combined use of radar and other sensors. During the early phases of such an engagement, detection will be intermittent and the system can be easily confused by false alarms. Based on these initial intermittent detections, additional sensor resources must be allocated to determine which of them represent valid targets. This additional resource may be in the form of different wave forms or sensing from other platforms. Initially, individual targets will be unresolvable and raid assessment may be necessary to determine how many targets are present. For low-flying targets, multipath interference can be a significant problem for radar systems. In this case, the high angular resolution of an imaging sensor can be particularly useful. Once targets have been detected and localized, they need to be classified and threatening targets intercepted. To classify them, sensors may switch modes to provide classification signatures. However, this is only useful if targets are sufficiently localized by the sensors for the classification modes to be effective. Finally, ownship maneuvers may be required during weapon engagement to achieve favorable launch or sensing geometry.

Current approaches to sensor management and data fusion suffer from a number of deficiencies when confronted with this type of coupled problem. First, most existing systems model targets as a collection of independent objects. Assuming object independence is equivalent to assuming that the joint probability density for the targets factorizes into a simple product form. Such a form fails to portray the correlation in the density that arises, for example, when targets are crossing or close together as in convoy. The existence of this correlation effect means that tracking can be improved for close targets if additional sensor resource is allocated to them. However, if the data fusion system fails to model the

correlation effect, then determining how to allocate sensors for close targets is more difficult.

Second, many existing sensor management schemes are based on relatively limited subsets of the information provided by the data fusion system, such as predicted position or velocity error. Such schemes have no way to evaluate the shared utility of non-commensurate quantities such as velocity sensitive sensor modes against modes that are primarily effective in classifying well-localized targets.

The third and final issue in most tracking systems is that they are based on Kalman filter estimation, which yields a Gaussian approximation of the probability density of the multitarget state. There are many situations where such an approximation breaks down. Poorly localized targets often have highly multimodal densities that are not well-approximated by a Gaussian form, or any other standard form for that matter. Geometry effects such as multipath interference can lead to multimodal densities. Nonlinear dynamics can lead to highly non-Gaussian densities, even if they remain mono-modal. Finally, the Kalman approach requires that it be possible to associate sensor measurements with corresponding objects in the state estimate, an action that can easily fail, especially at low signal-to-noise ratios (SNR). Techniques to relax the Gaussian assumption inherent in the Kalman approach, thereby treating the probability density of the target state in a more general manner, are referred to as nonlinear filtering methods.

This project addressed three fundamental issues required for effective sensor management, issues suggested by the tactical scenario above. First, a framework is needed to a) describe the disparate interacting components of the tactical scene, and b) to account for uncertain target dynamics and imperfect measurements. For example, during the early stages of this engagement, not only were the individual target locations uncertain, but the number of targets itself was uncertain partly because the targets were closely spaced. Later, the target number may be well-resolved but the target locations and classes are uncertain. For effective sensor management to take place, these joint uncertainties must be modeled by multitarget data fusion systems. Second, measures of expected utility for sensing alternatives are needed. The complexity of the joint multitarget uncertainty makes this task particularly difficult. Third, numerical methods are required to approximately solve the measurement-based estimation problem for our stochastic system, and to evaluate the resulting utility measures.

Setting the Stage

Bearing in mind the limited amount of formal development that had occurred in sensor management prior to the early 1990s, we took up the problems identified in the last section. To begin, the state of sensor management was assessed in the 1994 paper “Chasing the Elusive Sensor Manager” by Musick and Malhotra, [1]. This paper described the sensor management problem, reviewed its history, and sketched various techniques that might contribute to a rigorous and effective general theory to underpin our research. This paper was sobering for the large number of potentially viable techniques it found – the situation cried out for visionary insights or at least shrewd tactics. This paper has been referenced often by other researchers over the ensuing years.

MATHEMATICAL RESULTS

Discrimination Gain

In the months just prior to the start of this project, Keith Kastella and Wayne Schmaedeke, both then working at Unisys, had begun to look at sensor management techniques that used information metrics such as entropy and discrimination. Building on their preliminary work, over the period 1993-1997 this project developed and evaluated a technique to guide sensor allocation that employed the information measure called Kullback-Leibler discrimination. This technique used KL discrimination to predict how much the density function of the target state would shrink for any particular sensor mode and target/sensor pairing. The expected density shrinkage is a scalar that was termed the *discrimination gain*. The sensor management policy is then straightforward: search through all reasonable sensing combinations and select the one with the largest gain.

The idea of discrimination gain for sensor management originated at Unisys with Kastella and Schmaedeke. They developed this technique in 1994, and in 1997 Kastella published it in journal form [13]. Performance comparisons of discrimination gain with other plausible methods were conducted by Kastella, Schmaedeke and Musick, and published in 1994-96 at various conferences [2, 6, 7, 10].

The basic notion of discrimination gain for use in sensor management is quite simple and can be understood as follows. It assumes that a probability density is available to describe the state of a collection of targets in a region – call this the target state. The number of targets, their locations and target classes, and their motion condition may all be uncertain, with all of those uncertainties captured in the associated probability density. Furthermore, a sensor model is available to provide the measurement probability density given the target locations and classes. For any postulated measurement, the expected gain in discrimination of the updated target density with respect to the current density can be computed.

Suppose we want to assess the utility of a particular measurement in a particular region. Since the current density is available, the probability for alternative measurement outcomes (detection versus non-detection, say) can be computed, the density updated with the hypothetical measurements, and the resulting discrimination evaluated. Although computationally expensive, this computation can be carried out across all possible sensor allocations and the discrimination maximizing allocation can be selected.

To illustrate its power, reference [6] contrasts discrimination gain with three other methods for guiding a sensor that is searching for a single stationary dim target that occupies one cell in a large space of cells. This classic detection problem can be approached in many ways. The methods we investigated were named direct search, alert/confirm, and index rule. The direct search method, which allocates the same number of measurements to each cell in the search space, was chosen as a baseline because it is the simplest (and most naïve) policy possible. The alert/confirm method, which allocated additional measurements to any cell where a detection occurred on an initial look, was chosen because it

has been used in operational systems. Finally, the index rule was chosen because it is provably optimal for the special circumstances of this example (search for a single target in a case where the measurement density is complementary and symmetric).

A Monte Carlo simulation was used to investigate the performance of these four methods. A single simulation run consisted of 1000 measurements, with 1000 independent runs in the full ensemble of runs for each method. Each method was implemented to ingest measurements into the probabilistic target state using optimal Bayesian updating techniques. At intervals of 100 measurements, the simulation was paused and the decision-maker was forced to declare where it currently thought the target was located. For a particular measurement epoch, the percentage of wrong answers over 1000 repetitions of the experiment represents the probability of error. Figure 1 is a plot of that probability as a function of increasing sensing effort (gauged by measurements expended), by method.

Figure 1 shows that direct search performs worst and the index rule best [6]. Both results were expected. However, it was revealing and encouraging that discrimination gain performed almost as well as the optimal index rule, and considerably better than the operational method named alert/confirm.

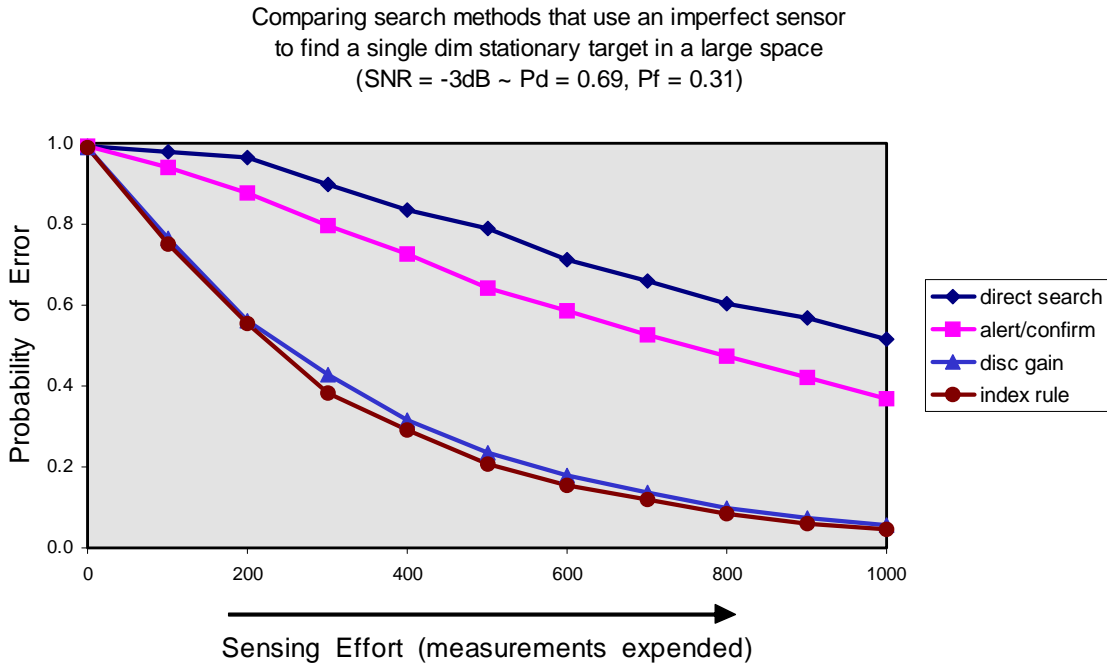


Figure 1. Comparison of four search methods

As matters turned out, this project would continue to use discrimination gain to guide sensor allocation throughout its history. Discrimination gain confers compelling advantages, the foremost ones being near optimality, the ability to work with generalized density functions, tractable computation (although relatively high computational burden),

and the ability to simultaneously balance the demands of diverse objectives like search, track and identification.

Joint Multitarget Probability

To treat the problem of simultaneous detection, tracking and identification in multitarget, multisensor applications, systems must model the joint uncertainty between all elements of a scene. This can be achieved through the use of the so-called *joint multitarget probability* (JMP). JMP is founded on Bayesian principles which permit using the standard tools of Bayesian analysis, including measurement updating via Bayes' rule, density propagation via the Fokker-Planck partial differential equation, and information theoretic notions such as discrimination. In particular, given the JMP framework, the expected gain in discrimination can be computed and used to guide sensor allocation decisions. This leads to a complete approach to data fusion based on a) tracking the probability for an unknown number of targets using a joint collection of multitarget probabilities, and b) maximizing the expected discrimination gain for each sensor dwell.

JMP was introduced by Kastella in [10] and may be understood as follows. JMP is based on the conditional probability density $p(x_1, \dots, x_n / Z)$ that a) there are exactly n targets in the scene, and b) they are located at x_1, \dots, x_n based on a set of observations Z . The joint collection of all such conditional probabilities for $n = 0, 1, \dots, N$ comprises a complete probability density, the JMP density, with a total mass that sums to one. Here N may be thought of as the maximum number of targets that could occur in the scene of interest. Given a measurement and a sensor model, Bayes' rule is used to update $p(x_1, \dots, x_n / Z)$ for all x_i and for each value of n . Target dynamics are modeled as Markovian, which leads to a time-evolution of $p(x_1, \dots, x_n / Z)$ that is independent of the measurement history. This Markov process can be modeled in discrete time or in the continuum limit, where the time-evolution of $p(x_1, \dots, x_n / Z)$ is then governed by the Fokker-Planck PDE determined by the dynamics of the individual targets. The expected discrimination gain for sampling a region with a sensor can be computed from $p(x_1, \dots, x_n / Z)$. The sensor is moded and directed to the region that maximizes the expected gain for each sample. In comparison to directly sampling all of the cells, optimizing the discrimination significantly increases the probability of detecting and localizing all of the targets.

Results obtained using JMP and discrimination gain in a variety of multitarget applications are reported in [10, 17, 19, 20, 21]. Figure 2 and Figure 3 illustrate such results for the problem of tracking two targets that pass one another while moving along a line in a one-dimensional space. Here the problem is to detect and track these targets in a region where the number of targets is not known a priori and where there is a very high false alarm rate (corresponding to about 0 dB SNR). In this test case the targets lie mostly between cells 4 and 8. Figure 2 shows the allocation of sensing effort through time using discrimination gain. Figure 3 shows the conditional probability that there are two targets in the space.

In Figure 2, location is indicated across the page, time goes into the page and the vertical axis gives average sensing effort. The targets are initially in cells 5 and 7. Between time 10 and 20, both targets are in cell 6. After time 20 the targets move apart. A simple sensor model is used where for each dwell the sensor can examine a single target cell. For each dwell the expected discrimination gain is computed, given the current value of the probability that one or more targets are in the cell. Then the cell with the highest expected gain is sampled. Ten sensor dwells are allocated for each time step. For the first time step the sensing effort is nearly uniformly distributed across the region. Once the targets are detected, discrimination gain automatically drives the sensor system to focus most of its effort in the target-containing region.

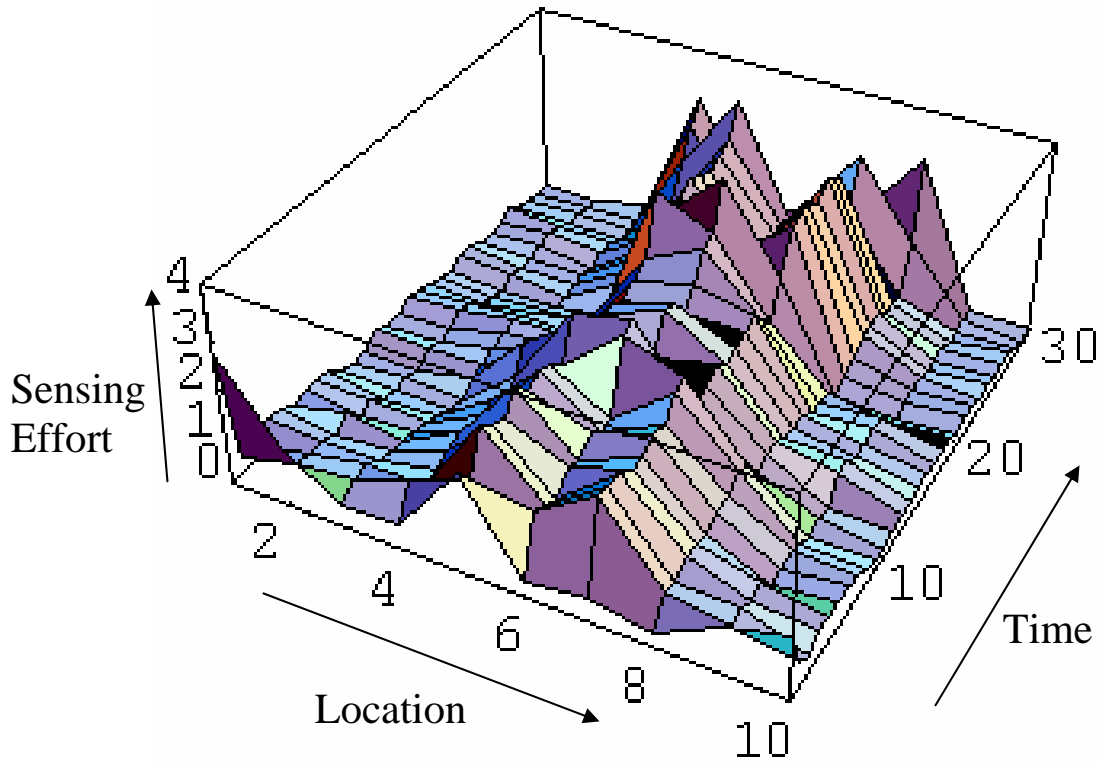


Figure 2. Sensing effort as a function of location and time for the 2-target example.

In Figure 3, the direct search (Dir) and discrimination gain (DG) methods are again compared, this time in terms of each method's ability to estimate the correct number of targets in the scene. The targets, which are moving, are co-located between times 10 and 20. During this period, the sensor cannot resolve them so they appear as one target and the probability for two targets falls. The upper curve was obtained using discrimination gain and the lower curve is the direct search result. Discrimination gain converges to 1 (the ideal answer) more quickly than direct search, representing improved performance.

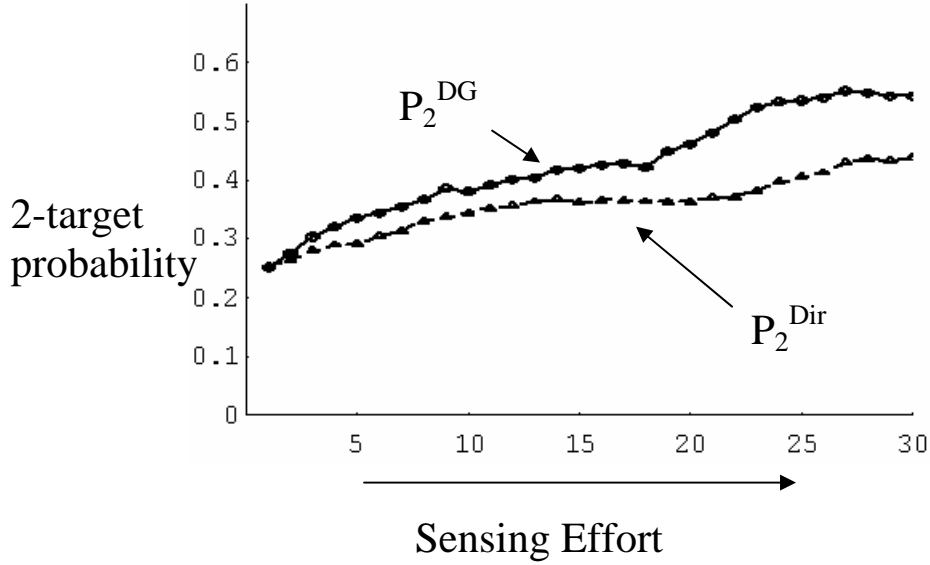


Figure 3. JMP probability P_2 that there are 2 targets in the scene.

JMP became the favored state representation for all efforts under Project SMFA. When target count is unknown, which is the usual case in practice, JMP provides significant advantages including mathematical rigor, a framework that can account for all sources of uncertainty, and compatibility with sensor management via discrimination gain. The biggest drawback with JMP is the high computational burden imposed in propagating the JMP density, especially when that solution is implemented by solving the Fokker-Planck PDE. Much of our work in the latter years of this project was directed at finding efficient means for solving this problem.

Nonlinear Filtering

The JMP formulation presents a classic problem in nonlinear filtering (NLF). The Bayesian foundations of NLF were laid in the context of single-target tracking and date to the 1960s. The feature that most distinguishes NLF from Kalman filtering and its many offspring is this: NLF uses a representation for the probability density of the target state that is entirely general, whereas Kalman assumes that density is Gaussian. Of course, when its assumptions hold, Kalman is the preferred approach because of its computational simplicity. The advantage of a general density is that it enables the nonlinear filter to treat the nonlinear effects of target dynamics and non-Gaussian measurements more realistically, thereby producing more accurate solutions. Furthermore, NLF has two particularly useful features: a) under the usual assumption of Markovian processes, the nonlinear filter is recursive, and b) the nonlinear filter is optimal within the Bayesian framework.

Early work on nonlinear estimation built upon the extended Kalman filter, leading to approximations such as Gaussian sum, point mass, and the unscented filter. These early approximations were pursued because the full nonlinear filter was generally viewed as unfeasible for real-time applications. Today, with faster computers and more efficient nu-

merical methods, NLF is a viable option for some applications. Today NLF techniques include spectral methods, separation of variables schemes, convolution methods, and Monte Carlo simulation schemes like particle filtering.

Tracking and identification problems are best modeled as having continuous-time target dynamics and discrete-time sensor measurements. After initialization, implementing a nonlinear filter consists of two basic steps: a) determining how the target state probability density evolves between measurements, and b) updating the target state probability density when a new measurement is obtained. The evolution of the target density can be determined by solving the Fokker-Planck (partial differential) equation (FPE), which describes how the target state density evolves between measurements under the influence of both deterministic and random effects. This entails solving a linear partial differential equation between sensor measurement epochs. The Bayes' rule implementation used for the measurement update is a relatively simple point-wise multiplication operation.

Most of the computational complexity and burden in NLF lies in propagating the target state density through time. If we assault the problem directly by employing finite difference methods to solve the FPE, an open question is which finite difference method is best in multitarget tracking and identification applications. To investigate this question, Kastella and Zatezalo developed and tested a variety of PDE solvers, including one based on the so-called Alternating Direction Implicit (ADI) finite difference scheme [11, 29]. ADI has a rigorous mathematical basis and its computational complexity is proportional to the number of grid nodes M used to approximate the target density, i.e. complexity is $O(M)$. Apparently, ADI's utility for problems of this type had not been previously recognized.

To illustrate ADI performance, consider the problem of detecting and tracking a dim target moving in a two-dimensional space where target dynamics are non-linear and image measurements are non-Gaussian. In particular, noise corrupts the image, producing an SNR of 3 dB. Additional problem facts include:

- the area of interest (AoI) is 6.4 km on a side;
- the maneuvering target travels in this AoI at 100 m/s for 70 sec, making a 1 G hairpin left turn over the sub-interval (20, 50) sec, see Figure 6;
- maneuvers are modeled as “nearly-coordinated”, with a stochastic motion model that contains five states, $X = [x, \dot{x}, y, \dot{y}, \omega]^T$;
- these five states are related nonlinearly, the deterministic part of the stochastic motion model being $f(X) = [\dot{x}, -\omega \dot{y}, \dot{y}, \omega \dot{x}, 0]^T$;
- the sensor is a downward-looking device that produces pixilated images of the entire AoI at 1 sec intervals;
- each sensor image is a 64×64 array of measurements, with pixels 100 m square;
- sensor intensity errors are distributed as Rayleigh noise;
- the filter is initialized with a uniform density over the full range of each variable.

Simulation results obtained for this situation are shown in Figure 4 through

Figure 6. Figure 4 shows a single image from the Rayleigh imaging sensor – note how difficult it is at 3 dB to tell which pixel contains the target.

Figure 6 shows six snapshots of the evolving marginal for the x-y target position. (Note that intensity values in these six marginals are plotted on a logarithmic (dB) scale.)

Figure 6 portrays an improving situation, e.g. the marginals are growing more compact and the error ellipses (not shown) contract by at least an order-of-magnitude during the 70 sec scenario. In most runs, about 20 sec (20 image scans) were required to localize the target in x-y position, 30 sec for x-y velocity, and 40 sec for ω . Once converged, the estimates are maintained through the remainder of the scenario.

Figure 6, an ensemble average of 10 runs, shows a well localized target through the straight portions of the trajectory but a significant increase in uncertainty during the 1 G turn itself. The turn rate ω has its greatest influence on the other states during the turn, and, as asserted above, is the state that was most difficult to estimate. These results are consistent with RMS error plots (not shown) across the ensemble. Raising SNR to 5 dB allows tracking through the turn to be substantially tighter.

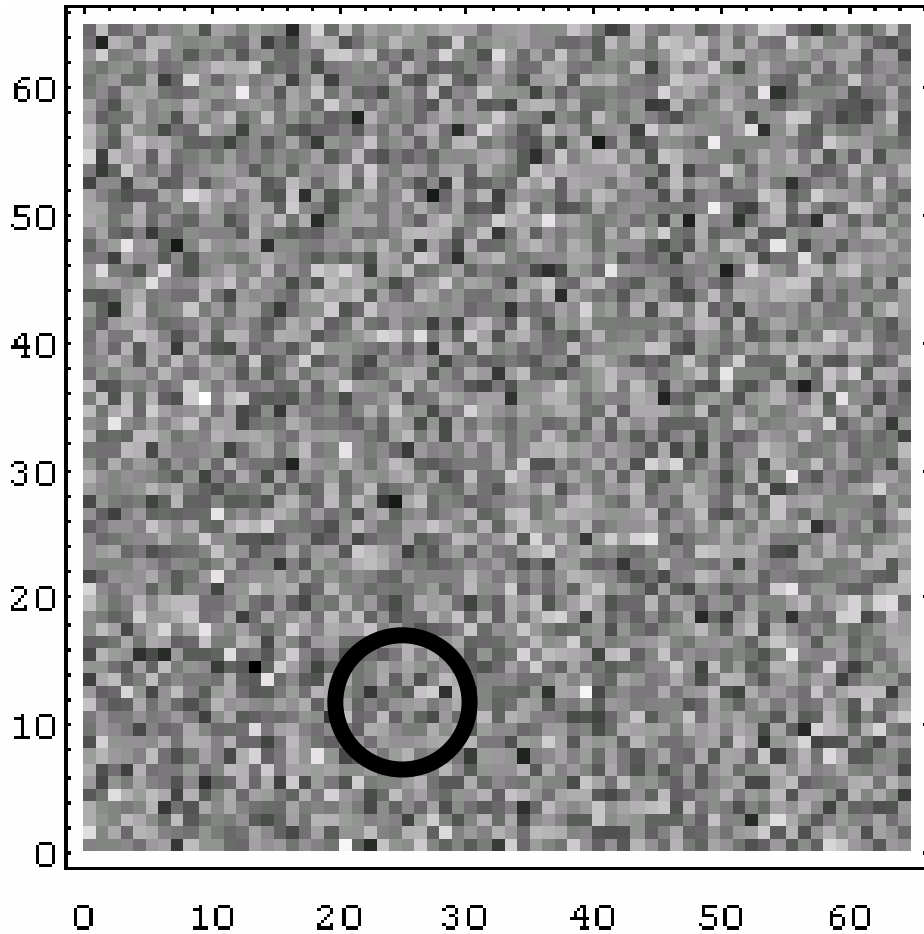


Figure 4. A single image at -3 dB, target at (27, 13)

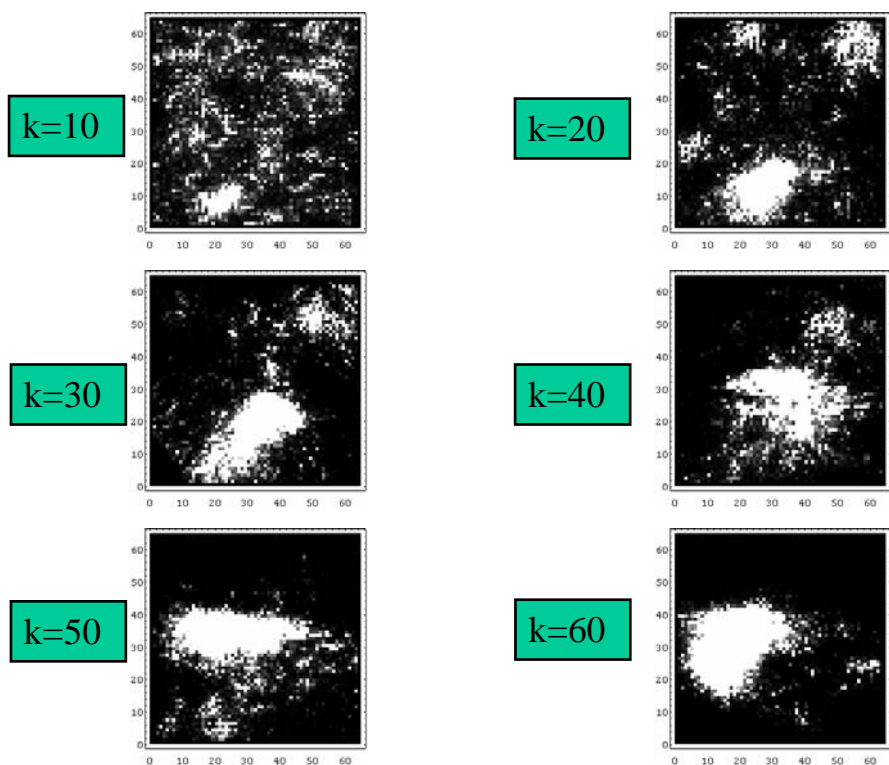


Figure 5. Position marginal, average over 10 runs

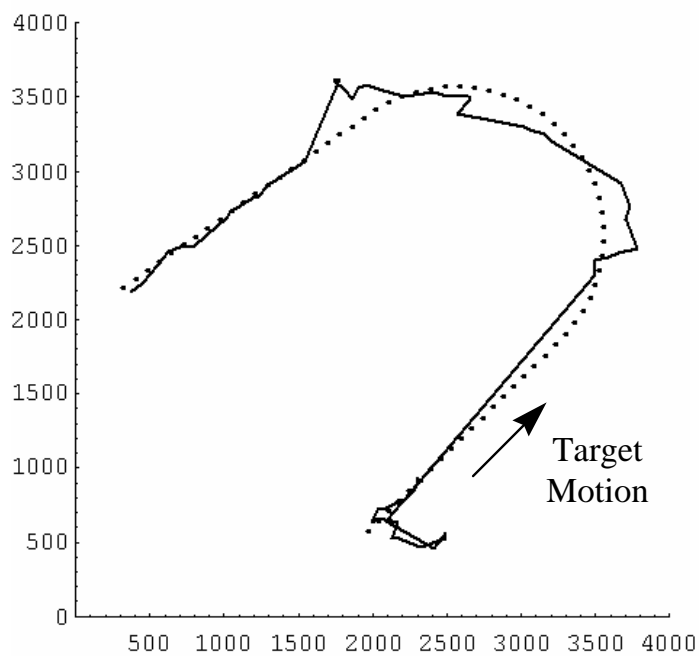


Figure 6. Low SNR image tracking, average over 10 runs, true dotted, estimate solid

TENET

In 1999, we began to study Monte Carlo methods for NLF as an alternative to directly solving the Fokker-Planck equation. Monte Carlo methods such as particle filtering emerged in the 1990s and quickly became prime candidates for the numerical solution of nonlinear problems. Such methods represent a probability density like JMP with a collection of particles that become dispersed over the probability density in numbers proportional to that density's mass concentrations. All particles are time-propagated per the system dynamic models via Monte Carlo simulation. At measurement update epochs, particles are evaluated by sampling the measurement at the discrete particle points and weighting the result according to a "proposal density". This step is called importance sampling. These resulting weights are used as the empirical sampling of the joint density of the state conditioned on the measurement. During the sampling step, particle filtering may generate many particles of low importance due to using randomization in the proposal process. A resampling step is used to replace low importance particles with higher importance particles so the particle distribution better represents the a posteriori density. This approach is fully Bayesian.

In 2000, Musick, Kastella, Kreucher and Greenewald developed a challenge problem in nonlinear filtering around a dim target tracking application. This challenge problem, which we named TENET (TEchniques for the Nonlinear Estimation of Tracks), was devised to encourage wider participation by the research community in NLF studies. TENET was introduced at a two-day workshop in February 2001 in Dayton that was hosted by AFRL/SNAT and attended by over 40 researchers, most of whom were active in NLF and/or in tracking. A web site was created at the following URL to facilitate distribution of the TENET software and documentation [43].

<https://www.vdl.afrl.af.mil/programs/tenet>

This is an open website, available to anyone who wishes to participate in the TENET NLF challenge problem.

References [38, 39, 41] are TENET-related conference papers written by contributors to Project SMFA. Although the TENET software and documentation have been downloaded some many times over the last five years, TENET has been cited only 14 times in related NLF papers over that same period. Furthermore, to our knowledge only one study has been conducted that used the TENET low-SNR scenario directly. Although one can never be sure about what motivates others, based on this low level of interest it seems clear to the authors of this report that capable and productive researchers are reluctant to undertake demanding work that has little prospect of financial return. Thus our failure to follow through with funding and other actions for this research effectively wasted the promising start that occurred in 2000-2001.

Applications of Nonlinear Filtering

This section describes several problems of Air Force interest that were addressed using SMFA technology.

Tracking in the presence of multipath interference

When radar observes a target near a reflecting surface such as the sea, it will generally receive echoes from both the target and the nearby surface. If the viewing geometry is constructive (as it is when the target is observed at low grazing angles over the surface), both echoes arrive at nearly the same time from nearly the same direction, creating reception patterns known as multipath interference. Such interference degrades detection performance and makes the direct echo from the target difficult to resolve from the reflecting surface echoes. Ultimately, multipath interference leads to difficulties in estimating target altitude above the surface. Although radar designers have found means in both hardware and signal processing to deal with radar multipath, current solutions are expensive and inaccurate, leaving much room for improvement.

In this study, NLF methods were used to exploit target motion to solve the altitude estimation problem. Ideally, target altitude could be estimated directly from the probability density of the radar measurement conditioned on target range and altitude. This direct approach is usually unfeasible because the measurement density generally has many false peaks that yield multiple solutions for target altitude. However, as target range varies, the locations of the false peaks fluctuate rapidly while the true peak steadily tracks target altitude.

In [20], Kastella and Zatezalo describe a nonlinear filter that exploits these measurement density peculiarities to estimate target altitude. This nonlinear filter recursively computes the probability density for altitude and altitude rate conditioned on the radar measurement sequence. The time evolution of this density between measurements is determined by the FPE, which is solved in real-time using the ADI finite difference scheme. The radar measurement density is computed from a physical model and used to update the conditional density of the target state using Bayes' rule.

In simulation testing with a typical shipboard radar that made measurements at 10 Hz, the nonlinear filter was able to reliably acquire and track transonic targets through mild maneuvers to produce an accuracy of about 12 m RMS (root-mean-square) in altitude, and 7 m/s RMS in altitude rate. These results demonstrate the feasibility of tracking in the presence of multipath interference using NLF techniques.

Association-free bias estimation

Nonlinear filtering research has consistently shown that by directly estimating the probability density of a target state using a track-before-detect scheme, weak and densely-spaced targets can be tracked, and data association can be avoided. Data association, which associates measurement reports with tracks, imposes a heavy burden on tracking, both in its design where complex data management structures are required, and in its execution which often levies a heavy computational burden. Therefore, avoiding data association can have significant advantages. However, a concern had long existed that data association is essential for estimating and correcting sensor biases, which are nearly always present.

This effort demonstrated that target tracks and sensor biases can be estimated simultaneously using association-free NLF methods based on the JMP representation. We began by defining a state consisting of target locations and a slowly drifting sensor bias. Stochastic models for state dynamics and for the measurement function were presented. A track-before-detect nonlinear filter was constructed to estimate the joint density of all state variables. A simulation that emulates estimator behavior was exercised under low SNR conditions. Simulation results showed that RMS values for both kinematic and bias states contracted as measurements were accumulated over time. This work, which is documented in [27, 30, 42], extended the useful range of NLF methods in tracking.

Tracking through radar clutter

The objective of this task was to track a single moving vehicle using measured radar data from a DARPA data collection. The technical challenges to achieving accurate estimation with this data were clutter that was intermittently heavy, data anomalies, and vehicle observations that changed radically in shape and size as the vehicle maneuvered over a variable ground terrain.

Several methods are available to track moving targets in clutter and noise from sensed kinematic and identity data. Among the most capable is track-before-detect (TBD), which delivers performance at lower ratios of signal-to-clutter-plus-noise (SCNR) than conventional tracking methods. Against isolated single-cell targets for example, TBD can detect and track at SCNRs as low as 0-6 dB.

This paper [44] explored the performance of TBD in scenarios involving multiple closely-spaced vehicles where radar sensors delivered a combination of kinematic and identity data. The identity data are range-profiles, obtained from a high range resolution (HRR) mode of the radar, that are used to help gauge the severity of vehicle maneuvers, while the kinematic data are ground moving target indications (GMTI). The TBD estimator, which is implemented using particle filter methods, is able to exploit the structure in the vehicle signature to better handle corruptors like poorly-modeled kinematics, clutter and noise. This paper described the TBD estimation method, discussed the experiments that were performed to test the method using real GMTI/HRR data, and presented the simulation and metrics that were used for evaluation. Results show that the method was able to operate at low SCNR in stressful estimation situations.

A method for finding distributed objects

Detecting and identifying distributed objects in an image is a recurrent problem in Automatic Target Recognition (ATR), and in application areas like astronomy, speech recognition, and biomedical imaging. Part of the challenge of such problems lies in the fact that the individual “spots” that comprise the distributed object may hold little intrinsic identification information. In such cases, identification can only be assured when the entire distributed object conforms to the expected pattern. In this work, knowledge of an object’s geometric shape and spot configuration makes its detection possible, even amid heavy clutter.

This paper [50] appeals to particle filtering methods to detect, localize, and identify a distributed object in a single cluttered image. By maximizing the joint probability that a particular collection of spots is the object of interest, the decision can be made with an acceptable error rate. The setting for this work is a government program that has restricted the release of information about the actual problem. Thus, the method is illustrated using a surrogate estimation problem that retains the essential attributes of the original problem. Results demonstrate that the proposed method yields acceptable error levels in both false detection and localization when the SCNR is above 5 decibels.

MACHINE LEARNING RESULTS

The sensor management problem provides several motivations for the investigation of machine learning (ML) techniques. First, tactical sensor managers will be required to operate under harsh time constraints – real-time optimization may not be feasible. ML could allow us to leverage off-line processing toward a complex on-line problem. Thus, it may offer an attractive computational tradeoff: extensive learning trials are traded for a quickly-computed, reactive policy (if state= x then action= u). Secondly, sensor management is plausibly modeled as a large, stochastic, Markov decision process (MDP). Such models can be optimally solved using dynamic programming (DP), but only when state propagation dynamics and objective functions are known to be linear and quadratic, respectively. In the absence of these conditions, an exact, closed-form solution cannot be found and, for reasonable size problems, the iterative DP approach becomes computationally prohibitive. ML allows us to closely approximate optimal but incalculable DP solutions while addressing the computational burden issues as well. Finally, sensor management appears to exhibit complex mathematical relationships between actions and consequences. A precise, closed-form expression for this has not been obtained – ML provides a means to learn to approximate this relationship. Ideally, ML obtains the action-consequence relationships in the mean sense (this is provably optimal for MDPs).

In carrying out our ML research, two distinct approaches were taken: Reinforcement Learning (RL) and Virtual Associative Networks (VANs). Both theoretical extensions and applications were explored. This work is briefly described in the next two subsections.

Reinforcement Learning in Sensor Management

The three points mentioned above provided the impetus to study Reinforcement Learning (RL) for sensor management. However, there are maturity issues with RL which hamper its effectiveness. Several questions in particular arise;

- How can we best set the learning (or synthesis) parameters in order to maximize our success?
- How can we judge the performance of a learned policy on-line and gracefully degrade in the face of changing conditions?
- How can we use RL in conjunction with other on-line techniques (i.e., discrimination gain)?

These questions needed answers in order to improve the plausibility of a reinforcement learning approach for sensor management.

In FY96, SMFA work concentrated on finding and applying analytical techniques that could help understand and predict the performance of RL in different environments. We made incremental improvements to the simulation model from FY95 (added more realistic object motion, a Linear-Quadratic-Gaussian case) and began to look at the performance of Temporal Difference Learning (TDL) for various cases (variations of system parameters). We observed that TDL is sensitive to synthesis parameters and in different ways for the various learning environments. The variation of performance for TDL was

significant: well over an order-of-magnitude difference in the root-mean-square error of the Value Function was observed in different conditions. This sensitivity of performance to operating conditions motivated us to find methods to predict the performance of RL in sensor management and other problems.

At this time we also began to examine various techniques for predicting the behavior of RL systems via simulation. The studied techniques included: 1) optimal stopping to predict the performance after a given, finite amount of training; 2) a body of work in sampling theory based on the central limit theorem was used to predict rate of convergence of the algorithms; and 3) an ordinary differential equation (ODE) method was used to predict performance asymptotes for infinite training times.

The average case behavior of RL in differing conditions can be studied by using the third method, the ODE method. Ljung's results in particular allow one to characterize all synthesis and system parameters in the ODE. One can then study the family of ODEs based upon the family of learning environments. Further, this could be applied to a linear quadratic Gaussian (LQG) problem so that the predicted asymptotes can be compared to an optimal (closed-form) solution. This method was successfully used to predict the performance of Temporal Difference Learning in various simple scenarios containing limited numbers of states and possible observations, as well as simple state transition laws. This use of the ODE Method was expanded to allow for predictions of performance on more complicated scenarios reflective of sensor management.

In FY97 we formulated new RL-directed search policies based on TDL. While synthesizing these algorithms we discovered several fundamental challenges for the application of RL to the static target detection problem. First, the challenge of posing the problem to the learning agent in a workable fashion was paramount. The continuous-valued hypotheses on which we learn to base current actions constitute infinite-dimensional state spaces. Learning over such spaces is a challenge for RL and generally requires the use of some function approximation methods (such as multi-layer perceptrons with back propagation). These, in turn, introduce a host of synthesis decisions and performance constraints which impact the amount of information an agent can process in a given time step (e.g., can the agent learn to simultaneously consider hypotheses from multiple cells/locations). Secondly, we encountered a sensitivity to our choice of incremental and final rewards for the learning agent. We examined information-theoretic incremental rewards based on entropy and cross-entropy as well as a formulation using true hypothesis error as a final reward. We found that the information-theoretic rewards produced a behavior which could only reach an asymptote at an error level of 0.08 (given more measurements the agent still identifies the target location wrong 8% of the time). This asymptotic performance was alleviated by training the agent with true hypothesis error as a final reward.

Using simulation to examine a detection scenario, we compared the performance of RL-directed search against an index policy that is optimal under certain narrow circumstances and an uninformed search policy which maintains a fixed search pattern regardless of new sensor information. RL-directed search proved best among these schemes: it per-

formed nearly as well as the optimal index policy when the narrow circumstances were obtained and much better when they were not (e.g., the case of multiple targets).

An overview of this work and our findings in Reinforcement Learning follows.

- Tested and documented the variation in performance of Temporal Difference Learning in [9] and in subsequent simulations.
- Studied statistical methods which can be used to predict and understand the performance of RL algorithms in sensor management problems modeled as Markov decision problems. Found several applicable methods: 1) Optimal stopping problem literature; 2) Central limit theorem/sampling theory literature; 3) Ordinary differential equation (ODE) literature. Pursued the ODE method by applying it to a simple Markov process to observe the method's ability to predict the asymptotic value of the value function. The ODE Method successfully predicted convergence values for these simple cases characterized by a limited number of states and possible observations, and simple state transition laws.
- Enhanced war game simulation to have more realistic object motion (based upon accelerations being applied and the laws of physics) and added an LQG scenario which allows for an optimal closed-form solution
- Compared TDL performance with different parameters (learning rates, eligibility horizons, etc) against each other and against the LQG solution.

Virtual Associative Networks for Sensor Management

In FY98, machine learning for sensor management refocused away from pure RL approaches and feed-forward neural networks toward Virtual Associative Networks. The impetus for this redirection came from limitations that were suspected early on [1] and then proven over the course of our investigations. Specifically, RL methods suffer when applied to problems of large scale. In the case of the sensor management problem, the large scale arises from the combinatorial explosion in both the state space and the decision space, a fact that necessitates excessively long training times and/or heuristic reductions in the number of states in the model. The scale of a realistic sensor management problem is simply so large that RL will always ultimately fail.

By contrast, VANs utilize a graph-theoretic representation of learned associations between features to drastically condense the decision space into a manageable size and form. The VAN paradigm is based upon experimental results in neuroscience which indicate that biological intelligence is rooted in mechanisms for association/dissociation across neuronal pools. The VAN paradigm instantiates this idea in the form of a self-partitioning, hierarchal graph structure whereby elementary features are represented as nodes which are connected by real-valued vertices representing associations between features. By weighting the vertex associations, subgraphs arise out of the structure that can be used to guide a search process. This shift to VANs was instigated by Mr. Malhotra in collaboration with Dr. Yan Yufik, the developer of the VAN paradigm.

In FY98 we applied VANs to the management of sensors in dense target environments where many objects must be scanned in a time-stressed situation. We assumed features

are extracted by the sensors with some noise and we use these features to recognize objects which have differing priorities. The goal is to recognize objects of interest (targets) within the large ensemble as quickly as possible. We compared various VAN-based strategies against a random search (as a benchmark) and achieved nearly two orders-of-magnitude increase in speed. This problem is intended to be broadly representative of the challenges associated with sensor management in reconnaissance missions.

In FY99 we continued to investigate how machine learning techniques based on VANs could be applied to problems involving sensor management. Our investigations were centered in three areas.

In the first we gathered information and refined our tool base to model the problem of managing the sensors in a geographically distributed reconnaissance scenario. Here we considered means for routing homogenous unmanned aerial vehicles (UAVs) in a dense target environment in which targets may dynamically appear or disappear in a probabilistic fashion. The task involves planning and re-planning UAV routing and sensor activity to maximize some measure of performance (probability of correct classification, expected target coverage, etc.). This situation, which we treated as a variant of the classic vehicle routing problem, was investigated along that line.

In the second area we expanded our simulation abilities to more closely reflect the problem area described above. We introduced the routing aspect into the distributed sensor management problem as well as unique stochastic characteristics such as random winds and service times, and variable travel times (which depend on travel direction as related to wind direction). We applied VANs to plan UAV routes and schedule sensor activity. We compared various VAN-based strategies with a greedy routing method coupled with random search. We observed between one and two orders-of-magnitude performance advantage in terms of time to classify high priority targets.

In the third area we explored issues relating to the maturation of the VAN paradigm. This included exploring various graph partitioning algorithms (a key step for VANs), and their efficiency and suitability for large graphs. We also considered the introduction of the concept of reinforcement into the VAN paradigm. This can produce a more rigorous paradigm which will not require domain-specific knowledge and ad hoc methods to leverage the information stored in the weighted graph of the VAN.

In FY00 we continued to investigate the applicability of the VAN paradigm to problems involving sensor management and dynamic routing of platforms with onboard sensors. Our efforts resulted in improvements in two areas, including expanding the simulation/model for multi-platform intelligence, surveillance, and reconnaissance (ISR) missions to include more realistic operational conditions, and enhancing the theoretical foundations for VANs by introducing the mechanism of *reinforcement* (a.k.a. *reinforcement learning*) into the packet formation process. Our FY00 activities are described below.

First we gathered information, met with product organizations, and refined our knowledge relating to sensor management of geographically distributed reconnaissance assets.

Here we considered the problem of routing multiple, homogeneous UAVs in a dense target environment in which sensor tasks may dynamically appear or disappear in a probabilistic fashion - this problem domain is known in the scientific literature as the *general vehicle routing problem*. We expanded our model to include tasks with geometric constraints requiring sensor platforms to view target areas from a particular direction; this mimics current operational reconnaissance and battle damage assessment requirements. Further, we introduced the target-to-sensor *clustering* problem in which we account for the fact that several targets may be viewed by a single sensor “footprint”. This introduces an algorithmic requirement to associate targets to sensor footprints. These enhancements were cited as desirable in our discussions with ISR product organizations and were included in our updated simulations. We applied VANs to plan/re-plan platform routes and schedule sensor activity. We compared various VAN-based strategies with a greedy routing method that used random search. We continued to observe between one and two orders-of-magnitude performance advantage in terms of time to classify high priority targets with the larger gains being observed as target density and target constraints increased.

Finally, we continued to explore issues relating to the maturation of the VAN paradigm. Here the primary thrust involved the incorporation of the notion of reinforcement to guide the formation of clusters within the graph. The concept of reinforcement allows one to weight associations between certain features (nodes in a VAN’s graph-like structure) more heavily than others based upon the observed significance of actions with outcomes. Although this application of reinforcement to the VAN paradigm was new, we believe it helped to mature this approach for large-scale resource allocation problems such as sensor management and dynamic route re-planning.

In FY 01, we continued to investigate the VAN approach and its applicability to the general vehicle routing problem. In particular, we investigated needed theoretical extensions for VANs that would improve tractability and performance. Our efforts focused on alternative mechanisms for introducing reinforcement into the VAN model, as well as incorporating related concepts from approximate dynamic programming, such as MDP models. We also investigated using information-theoretic measures (such as entropy) to guide the associative processes (packet formation and dissolution), a key component of the VAN model. These investigations produced notional concepts for maturing the VAN approach to sensor management.

In FY01 we also expanded our model to include precedence constraints between tasks and grouped tasks – this reflects current operations in which tasks may be ordered or grouped to accomplish specific objectives such as geo-locating a target, or maintaining identification of moving targets under difficult conditions. We applied VANs to plan/re-plan platform (UAV) routes and schedule sensor activity. Again, we compared various VAN-based strategies with a greedy routing method and observed roughly an order-of-magnitude performance advantage in time to identify all high priority targets with the larger gains being observed as target density and target constraints increased.

NEW DIRECTIONS

After seven years of effort, the mathematical line of attack in Project SMFA had developed mathematical and information-based foundations for sensor management in multi-target multisensor settings that were complete, theoretically rigorous, and high performing, even under difficult conditions such as crossing targets and low SNR. Although these accomplishments were significant, the computational burden for employing those foundations was extreme and our efforts to lower that burden (e.g. [19, 29]) had met with only limited success. Clearly, new ideas and more concerted efforts were needed if principled sensor management was to become reality.

In 2000, Dr. Kastella led a General Dynamics (GD) team that won an award for the DARPA program called Integrated Sensing and Processing (ISP). This program was the brainchild of Dr. Dennis Healey and Dr. Douglas Cochran, the latter becoming its program manager. ISP's goals were to foster research in sensor management and related sensor signal processing disciplines in order to enhance their theoretical foundations. With DARPA instructions to uncover new and fundamental insights, Dr. Kastella's team sought to build on results from Project SMFA to implement an innovative, principled and practical system that could be expected to work in realistic multitarget multisensor environments. GD's work on ISP has recently concluded. This section synthesizes that work.

GD's ISP team consisted primarily of Dr. Keith Kastella and Dr. Christopher Kreucher. Dr. Alfred Hero of the University of Michigan worked closely with the GD team with funding from a related DARPA MURI titled "Sequential Multi-Modality Target Detection and Classification Using Physics-Based Models". Dr. Kreucher was a primary ISP contributor, earning his Ph.D. under Dr. Hero in the topic "An Information-Based Approach to Sensor Resource Allocation". His dissertation was focused wholly on the ISP problem.

In [54], Kreucher and Kastella summarize GD-ISP progress as requiring the following three interrelated developments. (The following descriptions are slight alterations of their words, made only to adjust for the context of this report.)

- *Bayesian Multitarget Tracking.* First, GD-ISP constructed a high fidelity non-parametric probabilistic model that captures the uncertainty inherent in the multi-target tracking problem. This was done via the joint multitarget probability density (JMPD¹), which is a single entity that probabilistically describes the knowledge of the states (e.g., position and velocity in 2 dimensions plus identification) of each target as well as the number of targets. Due to the nature of the target tracking problem, it is essential to capture the correlations in uncertainty between the states of different targets as well as the coupling between the uncertainty about the number of targets and their individual states. The JMPD captures these

¹ JMP and JMPD are identical probabilistic representations. However, the numerical techniques developed in ISP for solving the associated NLF problem were quite different from those developed in SMFA. Of the two, ISP's is undoubtedly more capable and preferred.

couplings precisely as it makes no inherent factorization, independence, or parametric form assumptions about the density. Due to the high dimensionality and non-parametric nature of the density, advanced numerical methods are necessary to estimate the density in a computationally tractable manner. To this end, GD-ISP developed a novel multitarget particle filter with an adaptive sampling scheme that automatically factorizes the JMPD when permissible, and provides a measurement directed bias for target addition and removal. This filter allows recursive estimation of the JMPD in a Bayesian setting. A recent reference on this work is [51].

- *Information-based Sensor Resource Allocation.* Second, GD-ISP used the estimate of the JMPD to make (myopic) sensor resource allocation decisions. As was done in SMFA, GD-ISP took an information-based approach, where the fundamental paradigm is to make sensor tasking decisions that maximize the expected amount of information gained about the scenario, as measured by the. (The Rényi Divergence is also called α -divergence, the parameter α defined on $(0, 1)$ where KL discrimination is a special case of Rényi as α goes to 1.) This unifying metric allowed GD-ISP to automatically trade between sensor allocations that provide different types of information (e.g., actions that provide information about position versus actions that provide information about identification) without any ad hoc assumptions as to the relative utility of each. A recent reference on this work is [52].
- *Multistage Sensor Scheduling.* Third, GD-ISP took up the problem of extending the information-based sensor resource allocation paradigm to long-term (non-myopic) sensor scheduling. This extension allows the consideration of long-term information gaining capability when making decisions about current actions. This aspect is particularly important when the sensor has time-varying target response characteristics due to sensor motion, the behavior of the vehicles being tracked, or dynamic terrain features. GD-ISP developed two numerically efficient methods of approximating the long-term solution, as the exact solution is computationally intractable. The first is an information-directed search algorithm which focuses the Monte Carlo evaluations on action sequences that are most informative. The second is an approximate method of solving the Bellman equation which replaces the value-to-go with an easily computed function that approximates the long term value of the current action. A preliminary report is available in [53].

GD’s final ISP report [54] contains dozens of results, insights and conclusions, a collection that we cannot do justice to here. We choose three results that we trust will illustrate the power and potential of an information-based approach to sensor management.

Figure 7 is a snapshot of an area of interest containing three targets. This figure contrasts performance with and without sensor management, the left panel being the case with sensor management via Rényi Divergence, and the right panel the case without where periodic scan is used. (Periodic scan was previously called direct search.) Targets are marked

with an asterisk, the (x,y) covariance spread of the filter estimate is shown by the ellipses, and the grey scale at the right of each panel indicates the number of times each cell has been measured at this time step (the total number of measurement looks is identical in each case). In the case of periodic scan, an entire row constituting one twelfth of the region is scanned at each time step, starting at the bottom and proceeding to the top before repeating (cells scanned at this snapshot epoch are indicated by the white stripe). With sensor management, measurements are used only in areas that contain targets. Here is a direct quote from [54]: “Qualitatively, in the managed scenario measurements are focused in or near cells that the targets are in. Quantitatively, the covariance ellipses calculated by the filter show that performance is significantly better in the managed scenario.” These results are typical of what happens with and without sensor management.

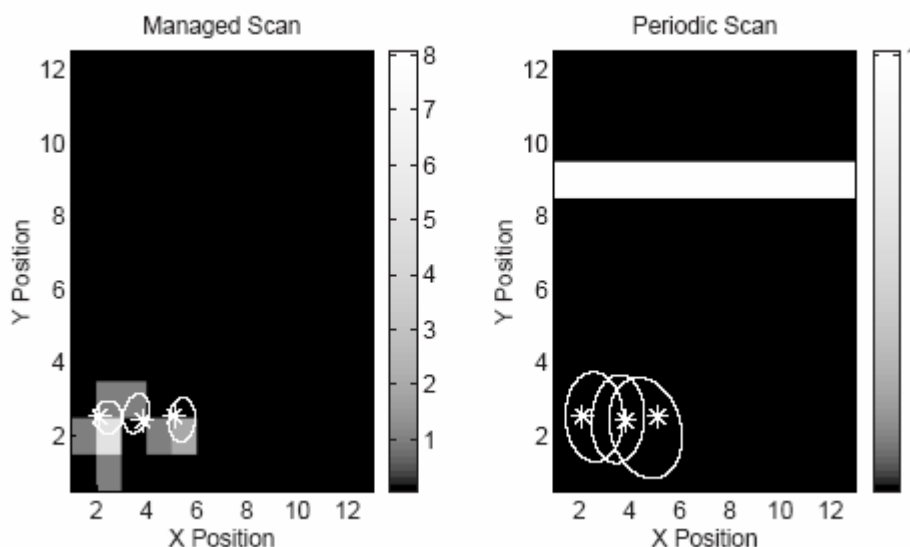


Figure 7. An illustration contrasting managed and non-managed tracking performance

Figure 8 illustrates the power of intelligent sensor management in terms of reducing sensing effort to achieve a particular goal. Again we are quoting from [54]. “A more detailed examination is provided in the Monte Carlo simulation results of Figure 8. We refer to each cell that is measured as a “look”, and are interested in empirically determining how many looks the non-managed algorithm requires to achieve the same performance as the managed algorithm at a fixed number of looks. The sensor management algorithm was run with 24 looks (i.e. was able to scan 24 cells at each time step) and is compared to the non-managed scheme with 24 to 312 looks. Here we take $\alpha = 0.99999$ to approximate the KL divergence. It is found that the non-managed scenario needs approximately 312 looks to equal the performance of the managed algorithm in terms of RMS error. Multi-target RMS position error is computed by computing the average RMS error across all targets. The sensor manager is approximately 13 times as efficient as compared to allocating the sensors without management. This efficiency implies that in an operational scenario target tracking could be done with an order of magnitude fewer sensor dwells.

Alternatively put, more targets could be tracked with the same number of total resources when this sensor management strategy is employed.”

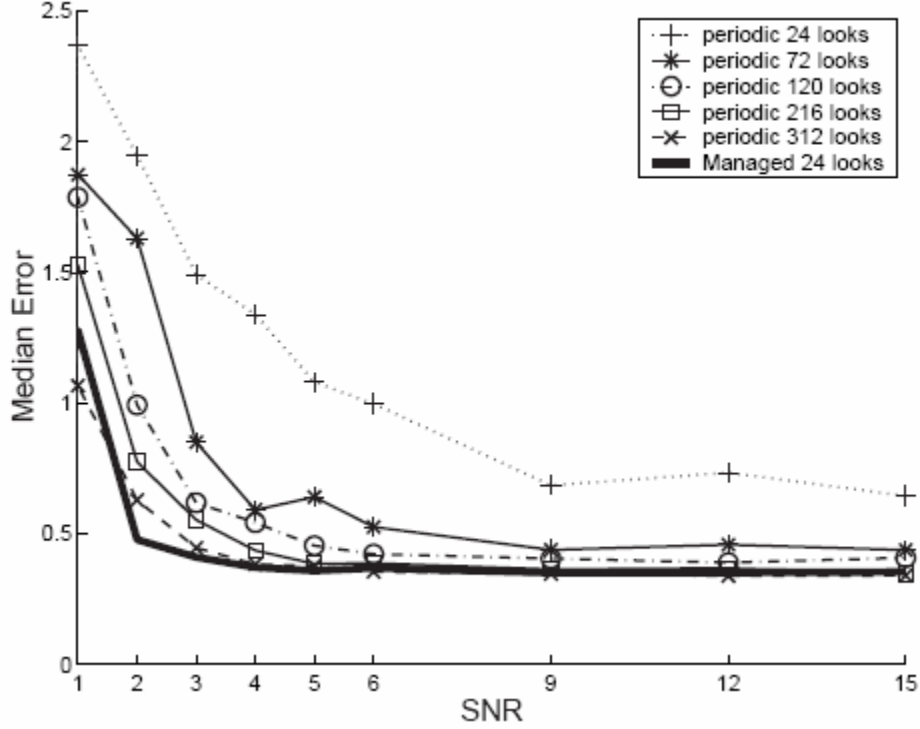


Figure 8. Quantitative comparison of intelligent and naïve sensor management.

Finally, Figure 9 presents empirical results illustrating the computational burden associated with using the GD-ISP algorithm in realistic scenarios. In particular, JMPD is implemented using 250 particles in the particle filter, sensor management is myopic with Rényi Divergence at $\alpha = 0.5$, and the measurements are thresholded. The simulation involves a 15×15 km ground surveillance region with moving targets numbering in the range 2 to 100. The imaging sensors are able to measure $100\text{m} \times 100\text{m}$ cells on the ground, meaning that at each time step there are 22,500 cells where the expected Rényi Divergence must be computed in order to determine the best sensing action. The simulation was implemented on an off-the-shelf 3 GHz Linux box.

We now quote from [54] again. “For equitable comparison in Figure 9, as the number of targets increases, the number of sensor resources increases (i.e. the number of sensing resources per target is kept constant throughout the algorithm). With modest optimization, a hybrid MatLab/C implementation of the algorithm is able to track on the order of 40 targets in real time and perform tracking and sensor management on 10 targets in real time.”

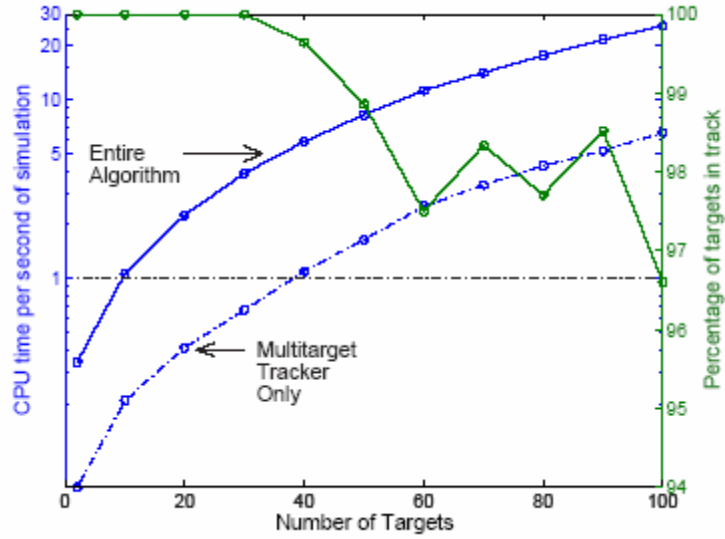


Figure 9. Execution time for myopic sensor management

In summary, GD-ISP produced the following advancements in the fields of target tracking and sensor management. Again, the words in the bulleted lists that follow are from [54] but adapted to this report.

- The development of a tractable particle filter based multitarget tracker to recursively estimate the joint multitarget probability density (JMPD). This approach simultaneously addresses estimation of target number and the state of each individual target, is nonparametric, and makes no assumptions of linearity or Gaussianity.
- The development of the Rényi Divergence metric for resource allocation in the multitarget tracking scenario. This method chooses sensor taskings in a manner that automatically trades between detection information, kinematic information, and identification information. The metric is general enough so that additional knowledge about the priority of each task can be incorporated.
- The extension of the information based sensor scheduling approach to multi-stage decision making through direct approximation and learning techniques.

As a result of GD-ISP work, we can draw the following broad conclusions about the problem domain and the overall utility of this project.

- By appropriate design of the importance density, it is possible to construct a tractable particle filter based multitarget tracker capable of estimating both the number of targets and the individual states of each in situations involving tens of targets and sensors.

- The Rényi Divergence framework for resource allocation is theoretically grounded and provides a natural method for trading the effects of different sensing actions.
 - The particle filter estimation and Rényi Divergence resource allocation algorithm are robust in the face of model mismatch.
 - Through marginalization and weighting, the Rényi Divergence can be used as a surrogate for task specific metrics.
 - In the case of discrete action spaces, this method provides a tractable method of resource allocation.
 - This method outperforms heuristic methods designed with domain knowledge.
- Multistage planning results in significant performance gain in situations where the system dynamics are changing rapidly.
 - Simple approximations to the MDP can provide good approximations to the multistage solution in many common scenarios.
 - Reinforcement learning methods are broadly applicable and can be used to address the multistage scheduling problem when training data and computational resources are available.

One final remark is in order. GD-ISP's demonstrated ability to manage tens of targets and sensors is a major improvement over the very small numbers that had typically been used in Project SMFA for testing and proving theory. Moreover, using advanced numerical techniques based on importance sampling, GD-ISP showed via simulation that tracking hundreds or thousands of targets is not necessarily computationally intractable (see Figure 9). More efficient implementations, for example through vectorization or parallelization, would certainly offer substantial gains that have not yet been explored. Thus, it is fair to conclude that GD-ISP reached a level of development where many diverse real-time applications in sensor management are now feasible for the first time.

REFERENCES

- [1] Stan Musick, Raj Malhotra, "Chasing the elusive sensor manager", *National Aerospace and Electronics Conference (NAECON)*, Dayton OH, pp 606-613, May 1994.
- [2] Wayne Schmaedeke, Keith Kastella, "Event-averaged maximum likelihood estimation and information-based sensor management", *SPIE Conference on Signal Processing, Sensor Fusion and Target Recognition III*, Orlando FL, pp 2232:91-96, April 1994.
- [3] Keith Kastella, "Mean-field theory and multitarget tracking", *SPIE Conference on Signal and Data Processing of Small Targets*, Orlando FL, pp 2235:388-393, April 1994.
- [4] Raj Malhotra, "Temporal Considerations in Sensor Management", *National Aerospace and Electronics Conference (NAECON)*, Dayton OH, May 1995.
- [5] Keith Kastella, "Event-averaged Maximum Likelihood Estimation and Mean-field Theory in Multitarget Tracking", *IEEE Transactions on Automatic Control*, Vol. 40, No 6, pp 1070-1074, June 1995.
- [6] Stan Musick, Keith Kastella, "Comparison of sensor management strategies for detection and classification", *9th National Symposium on Sensor Fusion*, Naval Postgraduate School, Monterey CA., March 1996.
- [7] Keith Kastella, Stan Musick, "The search for optimal sensor management", *SPIE International Symposium on Aerospace/Defense Sensing and Controls*, Orlando FL, April 1996.
- [8] Milton Cone, Stan Musick, "Scheduling for sensor management using genetic search", *National Aerospace and Electronics Conference (NAECON)*, Dayton OH, May 1996.
- [9] Raj Malhotra, "Convergence behavior of temporal difference learning", *National Aerospace and Electronics Conference (NAECON)*, Dayton OH, May 1996.
- [10] Keith Kastella, "Discrimination gain for sensor management in multitarget detection and tracking", Invited paper – *Tracking: Theory and Control Synthesis, IEEE-SMC and IMACS Multiconference CESA '96*, Lille France, July 1996.
- [11] Aleksandar Zatezalo, *Tracking and Detection for the Target State Model*, University of Minnesota M. S. Thesis, March 1997.
- [12] Raj Malhotra, "Achieving near-optimal sensor allocation policies through reinforcement learning", *Digital Avionics Systems Conference (DASC 96)*, Atlanta GA, November 1996.
- [13] Keith Kastella, "Discrimination gain to optimize detection and classification", *IEEE Transactions on Systems, Man and Cybernetics--Part A: Systems and Humans*, Vol. 27, No. 1, January 1997.
- [14] K. Kastella, M. Biscuso, W. Kober, J. K. Thomas, A. Wood, "Event-Averaged Maximum Likelihood Estimation Tracking for Fire-Control", *Proceedings of the 29th Southeastern Symposium on System Theory*, pp. 440-444, March 1997.
- [15] E. Blasch, R. Malhotra, S. Musick, "Using reinforcement learning for target search in sensor management", *IRIS National Symposium on Aerospace Sensor and Data Fusion*, MIT Lincoln Lab, Lexington MA, April 1997.
- [16] R. Washburn, A. Chao, D. Castanon, D. Bertsekas, R. Malhotra, "Stochastic dynamic programming for far-sighted sensor management", *IRIS National Symposium on Aerospace Sensor and Data Fusion*, MIT Lincoln Lab, Lexington MA, April 1997.
- [17] Keith Kastella, "Joint multitarget probabilities for detection and tracking", *SPIE AeroSense '97*, April 1997.
- [18] R. Malhotra, E. Blasch, J. Johnson, "Learning sensor-detection policies", *National Aerospace and Electronics Conference (NAECON)*, Dayton, OH, July 1997.
- [19] S. Musick, K. Kastella, R. Mahler, "A practical implementation of joint multitarget probabilities", *SPIE AeroSense '98*, Orlando FL, SPIE Vol. 3374, pp 26-37, April 1998.
- [20] Keith Kastella, Aleksandar Zatezalo, "Nonlinear filtering for tracking low elevation targets in multipath", *SPIE AeroSense '98*, SPIE Vol. 3373, Orlando FL, April 98.
- [21] Wayne Schmaedeke, Keith Kastella, "Discrimination gain and multiple model filters for sensor management in multitarget/multisensor tracking", *SPIE Aerosense '98*, Orlando FL, SPIE Vol. 3373, April 1998.
- [22] G. Sherman, D. Castanon, R. Malhotra, "Radar scheduling for the F-15E", *National Symposium on Aerospace Sensor and Data Fusion*, Atlanta GA, April 1998.

- [23] Yan Yufik, Raj Malhotra, "Virtual associative networks: A new paradigm for sensor fusion", *IEEE Joint International Symposium on Intelligence and Systems*, Rockville MD, May 1998.
- [24] Stan Musick, "Processing measurements of target count", unpublished technical memo, August 1998.
- [25] Raj P. Malhotra, Yan M. Yufik, "Virtual associative networks: A new paradigm for sensor fusion" *SPIE AeroSense '99*, SPIE Vol. 3719, Orlando FL, April 1999.
- [26] Raj P. Malhotra, Yan M. Yufik, "Graph-theoretic networks for holistic information fusion", *IEEE International Joint Conference on Neural Networks*, Washington D.C., July 1999.
- [27] Stan Musick, Keith Kastella, "Bias estimation in an association-free nonlinear filter", *Defense Applications of Signal Processing 99*, Starved Rock State Park, Utica IL, August 1999.
- [28] Raj P. Malhotra, Yan M. Yufik, "Information blending in Virtual Associative Networks: A new paradigm for sensor integration", *International Journal on Artificial Intelligence Tools*, Vol. 8, No. 3, pp. 275-290, Oct 1999.
- [29] Keith Kastella, "Finite Difference Methods for Nonlinear Filtering and Automatic Target Recognition", in *Multitarget Multisensor Tracking: Advanced Applications*, Vol. 3, Bar-Shalom and Blair (Editors), Artech House, 2000.
- [30] Stan Musick, Keith Kastella, "Track and bias estimation without data association," *Proceedings of SPIE AeroSense, Signal and Data Processing of Small Targets*, Orlando FL, SPIE Vol. 4052, pp258-269, April 2000.
- [31] K. Kastella, C. Kreucher, M. Pagels, "Nonlinear filtering for ground target applications," *SPIE AeroSense, Signal and Data Processing of Small Targets*, Orlando FL, April, 2000.
- [32] Raj P. Malhotra, "Virtual associative networks for reconnaissance in dense target environments" *Proceedings of Fourth International Conference on Cognitive and Neural Systems*, Boston University, Boston MA, May, 2000.
- [33] Keith Kastella, "A microdensity approach to multitarget tracking," *Proceedings of the 3rd International Conference on Information Fusion*, Paris, France, July 2000.
- [34] K. Kastella, B. Yeary, T. Zadra, R. Brouillard, E. Frangione, "Bias modeling and estimation for GMTI applications," *Proceedings of the 3rd International Conference on Information Fusion*, Paris, France, July 2000.
- [35] P. Shea, T. Zadra, D. Klammer, E. Frangione, R. Brouillard, K. Kastella, "Precision tracking of ground targets," *Proceedings of IEEE Aerospace 2000*, Big Sky, Montana.
- [36] Chris Kreucher, Keith Kastella, "Multiple model nonlinear filtering for low-signal ground target applications", *SPIE AeroSense, Signal and Data Processing of Small Targets*, Orlando FL, April 2001.
- [37] Stan Musick, Jeffery Layne, "Comparing distributed trackers", *Proceedings of the 4th Annual Conference on Information Fusion*, Montreal, Quebec, Canada, August 2001.
- [38] S. Musick, J. Greenewald, C. Kreucher and K. Kastella "Comparison of particle method and finite difference nonlinear filters for low SNR target tracking", *Proceedings of the 4th Annual Conference on Information Fusion*, Montreal, Quebec, Canada, August 2001.
- [39] John Greenewald, Stanton Musick, "A multiple-frame multiple-hypothesis method for tracking at low SNR", *SPIE Conference on Signal Processing, Sensor Fusion and Target Recognition XI*, Orlando FL, April 2002.
- [40] J. Hoffman, R. Mahler, B. Ravichandran, M. Huff, S. Musick, "Robust SAR ATR by hedging against uncertainty", *SPIE Conference on Signal Processing, Sensor Fusion and Target Recognition XI*, Orlando FL, April 2002.
- [41] Stan Musick, Keith Kastella, "Comparing numerical methods for nonlinear tracking applications", *Defense Applications of Signal Processing 01*, Barossa Valley Resort, Australia, June 2002.
- [42] Stan Musick, Keith Kastella, "Bias estimation in an association-free nonlinear filter", *Digital Signal Processing, A Review Journal*, Academic Press, ISSN: 1051-2004, August 2002.
- [43] Stan Musick, John Greenewald, "TENET simulator user's guide," unnumbered AFRL Report, July 2002.
- [44] John Greenewald, Stan Musick, "A method for employing features in nonlinear track filtering", *SPIE Conference on Signal Processing, Sensor Fusion, and Target Recognition XII*, Orlando FL, April 2003.
- [45] John Wegrzyn, Keith Kastella, "Intelligent sensor management to minimize detection error," *SPIE Conference on Signal Processing, Sensor Fusion, and Target Recognition XII*, Proceedings of SPIE Vol. 5096, pp. 502-513, April 2003.

- [46] C. Kreucher, K. Kastella, A. Hero, "Tracking multiple targets using a particle filter representation of the joint multitarget probability density", *Signal and Data Processing of Small Targets 2003*, Proceedings of SPIE Vol. 5204, pp. 258-269, San Diego, CA, August 2003.
- [47] Chris Kreucher, K. Kastella, A. Hero, "Information-based sensor management for multitarget tracking", *Signal and Data Processing of Small Targets 2003*, Proceedings of SPIE Vol. 5204, pp. 480-489, San Diego, CA, August 2003.
- [48] Keith Kastella, Richard Freeling, "Structured quantum search in NP-complete problems using the density of states", quant-ph/0109087 – submitted to *Physical Review Letters*, 2003.
- [49] S. Musick, J. Sherwood, T. Piatt, N. Carlson, "A simulation system for feature-aided tracking research", *Algorithms for Synthetic Aperture Radar Imagery XI, SPIE Defense & Security Symposium*, Orlando FL, April 2004.
- [50] J. Greenewald, S. Musick, T. Piatt, "A particle method for finding distributed objects", *7th International Conference on Information Fusion*, Stockholm Sweden, July 2004.

Some papers and a report by contributors to "Integrated Sensing and Processing":

- [51] C. Kreucher, K. Kastella, A. Hero, "Multitarget tracking using the joint multitarget probability density", *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 41, No. 4, October 2005.
- [52] C. Kreucher, K. Kastella, A. Hero, "Sensor management using an active sensing approach", *Signal Processing*, Vol. 85, No. 3, pp 607-24, 2005.
- [53] C. Kreucher, A. Hero, K. Kastella, D. Chang, "Efficient methods of non-myopic sensor management for multitarget tracking", *43rd IEEE Conference on Decision and Control*, Atlantis, Paradise Island, Bahamas, pp 722-27, Dec 2004.
- [54] C. Kreucher, K. Kastella, "Integrated Sensing and Processing (ISP)", General Dynamics Advanced Information Systems, Ypsilanti, MI, AFRL-SN-WP-TR-2006-1014, DTIC ADB317369, May 2006.